PORTAL USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

background processing task

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used: **background processing task** Found **92,556** of **213,097**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 Analysis of methods for scheduling low priority disk drive tasks

Eitan Bachmat, Jiri Schindler

June 2002 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '02**, Volume 30 Issue 1

Publisher: ACM Press

Full text available: pdf(161.76 KB) Additional Information: full citation, abstract, references, citings

This paper analyzes various algorithms for scheduling low priority disk drive tasks. The derived closed form solution is applicable to class of greedy algorithms that include a variety of background disk scanning applications. By paying close attention to many characteristics of modern disk drives, the analytical solutions achieve very high accuracy---the difference between the predicted response times and the measurements on two different disks is only 3% for all but one examined workload. This ...

2 On theories, methods and techniques: Core-task analysis in requirements gathering for an adaptive batch process control interface

Paula Savioja, Leena Norros, Leena Salo, Sauli Kivikunnas, Jari Karppinen

September 2005 **Proceedings of the 2005 annual conference on European association of cognitive ergonomics EACE '05**

Publisher: University of Athens

Full text available: pdf(363.65 KB) Additional Information: full citation, abstract, references

In this paper, we describe a research project in which Core-Task Analysis (CTA) is used to elicit functional requirements for an adaptive user interface of a batch process control system. The process under study is a continuous steel casting process the end products of which are steel slabs of varying size and quality. The paper is aimed to be methodological describing how technological requirements can be formulated with a work analysis method. Description of the prototype implementation shows ...

Keywords: adaptive user interface, continuous steel casting, core-task analysis, requirements analysis

3 Relevance and contributing information types of searched documents in task performance

Pertti Vakkari

July 2000 **Proceedings of the 23rd annual international ACM SIGIR conference on**

241499

92

STIC EIC 2100 Search Request Form

STIC Scientific & Technical Information Center

Today's Date:	What date would you like to use to limit the search?	
October 26, 2007	Priority Date:	Other:

Name Susan Ruggiero

AU 2167 Examiner # 77889

Room # 3C-05 Phone x1675

Serial # 10/763,081

Format for Search Results (Circle One):

PAPER DISK EMAIL

Where have you searched so far?

USP DWPI EPO JPO ACM IBM TDB

IEEE INSPEC SPI Other _____

Is this a "Fast & Focused" Search Request? (Circle One) YES NO

A "Fast & Focused" Search is completed in 2-3 hours (maximum). The search must be on a very specific topic and meet certain criteria. The criteria are posted in EIC2100 and on the EIC2100 NPL Web Page at http://ptoweb/patents/stic/stic-tc2100.htm.

What is the topic, novelty, motivation, utility, or other specific details defining the desired focus of this search? Please include the concepts, synonyms, keywords, acronyms, definitions, strategies, and anything else that helps to describe the topic. Please attach a copy of the abstract, background, brief summary, pertinent claims and any citations of relevant art you have found.

Is this request for a BOARD of APPEALS case? (Circle One) YES NO

Is this case a SPECIAL CASE? (Circle One) YES NO

Title: Redirecting a Response to a File Request ... to a mail server

- Request File
- if predetermined time has expired (from when request sent) and file requested is received → send file to target address
- if predetermined time has not expired since request sent and upon receipt of file → display file

background / batch (processing / task) search

+ time out

target address / e-mail address | timeout expiration / length of time predetermined

Inventor Chakravarty, Vijaylaxmi

IBM

STIC Searcher	Lucy Park	Phone	28667
Date picked up	10-26-07	Date Completed	10-26-07

Search and Information Resources Administration

ACM PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
background processing task and search and timeout
THE ACM DIGITAL LIBRARY
Feedback Report a problem Satisfaction survey

Terms used:
background processing task and **search** and **timeout**

Found **115,993** of **213,097**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 Smalltalk-80: the language and its implementation
Adele Goldberg, David Robson
January 1983 Book
Publisher: Addison-Wesley Longman Publishing Co., Inc.
Full text available: pdf(33.56 MB) Additional Information: full citation, abstract, cited by, index terms, review

From the Preface (See Front Matter for full Preface)

Advances in the design and production of computer hardware have brought many more people into direct contact with computers. Similar advances in the design and production of computer software are required in order that this increased contact be as rewarding as possible. The Smalltalk-80 system is a result of a decade of research into creating computer software that is appropriate for producing highly functional and interactive ...

2 System-level power optimization: techniques and tools
Luca Benini, Giovanni de Micheli
April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 5 Issue 2
Publisher: ACM Press
Full text available: pdf(385.22 KB) Additional Information: full citation, abstract, references, citings, index terms

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consu me energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

3 ObjectGlobe: Ubiquitous query processing on the Internet
R. Braumandl, M. Keidl, A. Kemper, D. Kossmann, A. Kreutz, S. Seltzsam, K. Stocker
August 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 10 Issue 1
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(251.44 KB) Additional Information: full citation, abstract, citings, index terms

We present the design of ObjectGlobe, a distributed and open query processor for Internet data sources. Today, data is published on the Internet via Web servers which

ACM PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
background (processing or task) and search and timeout
THE ACM DIGITAL LIBRARY
Feedback Report a problem Satisfaction survey

Terms used:
background processing or **task** and **search** and **timeout**

Found **115,993** of **213,097**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 Smalltalk-80: the language and its implementation
Adele Goldberg, David Robson
January 1983 Book
Publisher: Addison-Wesley Longman Publishing Co., Inc.
Full text available: pdf(33.56 MB) Additional Information: full citation, abstract, cited by, index terms, review

From the Preface (See Front Matter for full Preface)

Advances in the design and production of computer hardware have brought many more people into direct contact with computers. Similar advances in the design and production of computer software are required in order that this increased contact be as rewarding as possible. The Smalltalk-80 system is a result of a decade of research into creating computer software that is appropriate for producing highly functional and interactive ...

2 System-level power optimization: techniques and tools
Luca Benini, Giovanni de Micheli
April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 5 Issue 2
Publisher: ACM Press
Full text available: pdf(385.22 KB) Additional Information: full citation, abstract, references, citings, index terms

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consu me energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

3 ObjectGlobe: Ubiquitous query processing on the Internet
R. Braumandl, M. Keidl, A. Kemper, D. Kossmann, A. Kreutz, S. Seltzsam, K. Stocker
August 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 10 Issue 1
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(251.44 KB) Additional Information: full citation, abstract, citings, index terms

We present the design of ObjectGlobe, a distributed and open query processor for Internet data sources. Today, data is published on the Internet via Web servers which

PORTAL USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

search and timeout

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used: **search** and **timeout** Found **46,043** of **213,097**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 Using query mediators for distributed searching in federated digital libraries

Naomi Dushay, James C. French, Carl Lagoze

August 1999 **Proceedings of the fourth ACM conference on Digital libraries DL '99**

Publisher: ACM Press

Full text available: pdf(99.29 KB) Additional Information: full citation, references, citings, index terms

Keywords: distributed searching, query routing, search engine performance

2 Novel search environments: Integration of simultaneous searching and reference linking across bibliographic resources on the web

William H. Mischo, Thomas G. Habing, Timothy W. Cole

July 2002 **Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries JCDL '02**

Publisher: ACM Press

Full text available: pdf(279.99 KB) Additional Information: full citation, abstract, references, citings, index terms

Libraries and information providers are actively developing customized portals and gateway software designed to integrate secondary information resources such as A & I services, online catalogs, and publishers full-text repositories. This paper reports on a project carried out at the Grainger Engineering Library at the University of Illinois at Urbana-Champaign to provide web-based asynchronous simultaneous searching of multiple secondary information resources and integrated reference linking be ...

Keywords: asynchronous simultaneous search, digital library, openURL, reference linking, searchbots

3 Power reduction at circuit level: Determining the optimal timeout values for a power-managed system based on the theory of Markovian processes: offline and online algorithms

Peng Rong, Massoud Pedram

March 2006 **Proceedings of the conference on Design, automation and test in Europe: Proceedings DATE '06**

Publisher: European Design and Automation Association

PORTAL USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

search and timeout and "backgroud task"

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used: **search** and **timeout** and **backgroud task** Found **3,528** of **213,097**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 Using query mediators for distributed searching in federated digital libraries

Naomi Dushay, James C. French, Carl Lagoze

August 1999 **Proceedings of the fourth ACM conference on Digital libraries DL '99**

Publisher: ACM Press

Full text available: pdf(99.29 KB) Additional Information: full citation, references, citings, index terms

Keywords: distributed searching, query routing, search engine performance

2 Algorithms for multihop networks I: Revisiting the TTL-based controlled flooding search: optimality and randomization

Nicholas Chang, Mingyan Liu

September 2004 **Proceedings of the 10th annual international conference on Mobile computing and networking MobiCom '04**

Publisher: ACM Press

Full text available: pdf(371.21 KB) Additional Information: full citation, abstract, references, citings, index terms, review

In this paper we consider the problem of searching for a node or an object (i.e., piece of data, file, etc.) in a large network. Applications of this problem include searching for a destination node in a mobile ad hoc network, querying for a piece of desired data in a wireless sensor network, and searching for a shared file in an unstructured peer-to-peer network. We limit our attention in this study to the class of controlled flooding search strategies where query/search packets are broadcast a ...

Keywords: TTL, best worst-case performance, competitive ratio, controlled flooding search, data query, dynamic programming, file search, fixed strategy, optimal search, randomized strategy, strategy, wireless sensor networks

3 Efficient software model checking of data structure properties

Paul T. Darga, Chandrasekhar Boyapati

October 2006 **ACM SIGPLAN Notices , Proceedings of the 21st annual ACM SIGPLAN conference on Object-oriented programming systems, languages, and applications OOPSLA '06**, Volume 41 Issue 10

Publisher: ACM Press

PORTAL
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

"search engine" and (timeout or "length of time") and "backg

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used: **search engine** and **timeout** or **length of time** and **backgroud task**

Found **352** of **213,097**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Cluster-based scalable network services

Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: pdf(2.42 MB) Additional Information: full citation, references, citings, index terms

2 Quantifying and Improving the Availability of High-Performance Cluster-Based Internet Services

Kiran Nagaraja, Neeraj Krishnan, Ricardo Bianchini, Richard P. Martin, Thu D. Nguyen

November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing SC '03**

Publisher: IEEE Computer Society

Full text available: pdf(306.01 KB) Additional Information: full citation, abstract

Cluster-based servers can substantially increase performance when nodes cooperate to globally manage resources. However, in this paper we show that cooperation results in a substantial availability loss, in the absence of high-availability mechanisms. Specifically, we show that a sophisticated cluster-based Web server, which gains a factor of 3 in performance through cooperation, increases service unavailability by a factor of 10 over a non-cooperative version. We then show how to augment this W ...

3 DDoS defense by offense

Michael Walfish, Mythili Vutukuru, Hari Balakrishnan, David Karger, Scott Shenker

August 2006 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '06**, Volume 36 Issue 4

Publisher: ACM Press

Full text available: pdf(334.96 KB) Additional Information: full citation, abstract, references, index terms

This paper presents the design, implementation, analysis, and experimental evaluation of *speak-up*, a defense against *application-level* distributed denial-of-service (DDoS), in which attackers cripple a server by sending legitimate-looking requests that consume computational resources (*e.g.*, CPU cycles, disk). With speak-up, a victimized server encourages all clients, resources permitting, *to automatically send higher volumes of*

traffic. We suppose that attackers are a ...

Keywords: DoS attack, bandwidth, currency

4 A case study of automatically creating test suites from web application field data

Sara Sprenkle, Emily Gibson, Sreedevi Sampath, Lori Pollock
July 2006 **Proceedings of the 2006 workshop on Testing, analysis, and verification of web services and applications TAV-WEB '06**
Publisher: ACM Press
Full text available: pdf(184.56 KB) Additional Information: full citation, abstract, references, index terms

Creating effective test cases is a difficult problem, especially for web applications. To comprehensively test a web application's functionality, test cases must test complex application state dependencies and concurrent user interactions. Rather than creating test cases manually or from a static model, field data provides an inexpensive alternative to creating such sophisticated test cases. An existing approach to using field data in testing web applications is user-session-based testing. Previ ...

5 Scalability in MMOGs: Load balancing for massively multiplayer online games

Fengyun Lu, Simon Parkin, Graham Morgan
October 2006 **Proceedings of 5th ACM SIGCOMM workshop on Network and system support for games NetGames '06**
Publisher: ACM Press
Full text available: pdf(544.67 KB) Additional Information: full citation, abstract, references, index terms

Supporting thousands, possibly hundreds of thousands, of players is a requirement that must be satisfied when delivering server based online gaming as a commercial concern. Such a requirement may be satisfied by utilising the cumulative processing resources afforded by a cluster of servers. Clustering of servers allow great flexibility, as the game provider may add servers to satisfy an increase in processing demands, more players, or remove servers for routine maintenance or upgrading. If ca ...

6 Learning search engine specific query transformations for question answering

Eugene Agichtein, Steve Lawrence, Luis Gravano
April 2001 **Proceedings of the 10th international conference on World Wide Web WWW '01**
Publisher: ACM Press
Full text available: pdf(205.68 KB) Additional Information: full citation, references, citings, index terms

Keywords: information retrieval, query expansion, question answering, web search

7 Polling systems with server timeouts and their application to token passing networks

Edmundo de Souza e Silva, H. Richard Gail, Richard R. Muntz
October 1995 **IEEE/ACM Transactions on Networking (TON)**, Volume 3 Issue 5
Publisher: IEEE Press
Full text available: pdf(1.61 MB) Additional Information: full citation, references, citings, index terms

8 Using Time Instead of Timeout for Fault-Tolerant Distributed Systems.

Leslie Lamport
April 1984 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 6 Issue 2

Publisher: ACM Press
Full text available: pdf(1.81 MB) Additional Information: full citation, references, citings, index terms

9 Usage analysis: Semantic similarity between search engine queries using temporal correlation
Steve Chien, Nicole Immorlica
May 2005 **Proceedings of the 14th international conference on World Wide Web WWW '05**
Publisher: ACM Press
Full text available: pdf(238.57 KB) Additional Information: full citation, abstract, references, citings, index terms

We investigate the idea of finding semantically related search engine queries based on their *temporal correlation*; in other words, we infer that two queries are related if their popularities behave similarly over time. To this end, we first define a new measure of the temporal correlation of two queries based on the correlation coefficient of their frequency functions. We then conduct extensive experiments using our measure on two massive query streams from the MSN search engine, revealin ...

Keywords: query stream analysis, search engines, semantic similarity among queries

10 Search engineering 1: What's new on the web?: the evolution of the web from a search engine perspective
Alexandros Ntoulas, Junghoo Cho, Christopher Olston
May 2004 **Proceedings of the 13th international conference on World Wide Web WWW '04**
Publisher: ACM Press
Full text available: pdf(502.41 KB) Additional Information: full citation, abstract, references, citings, index terms

We seek to gain improved insight into how Web search engines shouldcope with the evolving Web, in an attempt to provide users with themost up-to-date results possible. For this purpose we collectedweekly snapshots of some 150 Web sites over the course of one year,and measured the evolution of content and link structure. Our measurements focus on aspects of potential interest to search engine designers: the evolution of link structure over time, the rate ofcreation of new pages and new distinct c ...

Keywords: change prediction, degree of change, link structure evolution, rate of change, search engines, web characterization, web evolution, web pages

11 Experiences with selecting search engines using metasearch
Daniel Dreilinger, Adele E. Howe
July 1997 **ACM Transactions on Information Systems (TOIS)**, Volume 15 Issue 3
Publisher: ACM Press
Full text available: pdf(428.65 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Search engines are among the most useful and high-profile resources on the Internet. The problem of finding information on the Internet has been replaced with the problem of knowing where search engines are, what they are designed to retrieve, and how to use them. This article describes and evaluates SavvySearch, a metasearch engine designed to intelligently select and interface with multiple remote search engines. The primary metasearch issue examined is the importance of carefully selecti ...

Keywords: WWW, information retrieval, machine learning, search engine

12 Education and evaluation: Implementation and evaluation of a quality-based search engine

Thomas Mandl

August 2006 **Proceedings of the seventeenth conference on Hypertext and hypermedia HYPERTEXT '06**

Publisher: ACM Press

Full text available: pdf(353.91 KB) Additional Information: full citation, abstract, references, citings, index terms

In this paper, an approach for the implementation of a quality-based Web search engine is proposed. Quality retrieval is introduced and an overview on previous efforts to implement such a service is given. Machine learning approaches are identified as the most promising methods to determine the quality of Web pages. Features for the most appropriate characterization of Web pages are determined. A quality model is developed based on human judgments. This model is integrated into a meta search eng ...

Keywords: quality models, quality search, web design, web metrics

13 Interaction: Investigating the querying and browsing behavior of advanced search engine users

Ryen W. White, Dan Morris

July 2007 **Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '07**

Publisher: ACM Press

Full text available: pdf(316.63 KB) Additional Information: full citation, abstract, references, index terms

One way to help all users of commercial Web search engines be more successful in their searches is to better understand what those users with greater search expertise are doing, and use this knowledge to benefit everyone. In this paper we study the interaction logs of advanced search engine users (and those not so advanced) to better understand how these user groups search. The results show that there are marked differences in the queries, result clicks, post-query browsing, and search succes ...

Keywords: advanced search features, expert searching, query formulation, query syntax

14 IDS, content filtering, Java, etc.: An architecture of pattern-oriented distributed meta-search engine

Tieli Sun, Wei Zhao, Zhiyan Zhao

November 2004 **Proceedings of the 3rd international conference on Information security InfoSecu '04**

Publisher: ACM Press

Full text available: pdf(348.40 KB) Additional Information: full citation, abstract, references

Normally, the implementation of meta-search engines adopts centralized methodology at present. When treating with massive data, the limitation of processing abilities of computers becomes a bottleneck of processing search results and increasing search speed. Existing search engines ubiquitously have many kinds of problems, such as the scope of searching is not wide enough, accessed sites are unbalance, primal content is scarce, and the information is not uniform. In contrast, distributed meta-se ...

Keywords: distributed, meta-search engine, pattern

15 Controlling the robots of Web search engines

J. Talim, Z. Liu, Ph. Nain, E. G. Coffman

June 2001 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '01**, Volume 29 Issue 1

Publisher: ACM Press

Full text available: pdf(686.42 KB) Additional Information: full citation, abstract, references, citings, index terms

Robots are deployed by a Web search engine for collecting information from different Web servers in order to maintain the currency of its data base of Web pages. In this paper, we investigate the number of robots to be used by a search engine so as to maximize the currency of the data base without putting an unnecessary load on the network. We adopt a finite-buffer queueing model to represent the system. The arrivals to the queueing system are Web pages brought by the robots; service corresponds ...

Keywords: Markov decision process, Web search engines, queues, web robots

16 Optimizing data aggregation for cluster-based internet services

Lingkun Chu, Hong Tang, Tao Yang, Kai Shen

June 2003 **ACM SIGPLAN Notices , Proceedings of the ninth ACM SIGPLAN symposium on Principles and practice of parallel programming PPoPP '03**, Volume 38 Issue 10

Publisher: ACM Press

Full text available: pdf(275.38 KB) Additional Information: full citation, abstract, references, index terms, review

Large-scale cluster-based Internet services often host partitioned datasets to provide incremental scalability. The aggregation of results produced from multiple partitions is a fundamental building block for the delivery of these services. This paper presents the design and implementation of a programming primitive -- Data Aggregation Call (DAC) -- to exploit partition parallelism for cluster-based Internet services. A DAC request specifies a local processing operator and a global reduction ope ...

Keywords: cluster-based network services, fault tolerance, load-adaptive tree formation, response time, scalable data aggregation, throughput

17 Attacks and cryptanalysis: Puppetnets: misusing web browsers as a distributed attack infrastructure

V. T. Lam, S. Antonatos, P. Akritidis, K. G. Anagnostakis

October 2006 **Proceedings of the 13th ACM conference on Computer and communications security CCS '06**

Publisher: ACM Press

Full text available: pdf(871.35 KB) Additional Information: full citation, abstract, references, index terms

Most of the recent work on Web security focuses on preventing attacks that *directly* harm the browser's host machine and user. In this paper we attempt to quantify the threat of browsers being *indirectly* misused for attacking third parties. Specifically, we look at how the existing Web infrastructure (e.g., the languages, protocols, and security policies) can be exploited by malicious Web sites to remotely instruct browsers to orchestrate actions including denial of service attacks, ...

Keywords: distributed attacks, malicious software, web security

18 Learning to find answers to questions on the Web

Eugene Agichtein, Steve Lawrence, Luis Gravano

May 2004 **ACM Transactions on Internet Technology (TOIT)**, Volume 4 Issue 2

Publisher: ACM Press

Full text available: pdf(4.49 MB) Additional Information: full citation, abstract, references, citings, index terms

We introduce a method for learning to find documents on the Web that contain answers to a given natural language question. In our approach, questions are transformed into new queries aimed at maximizing the probability of retrieving answers from existing information retrieval systems. The method involves automatically learning phrase features for classifying questions into different types, automatically generating candidate query transformations from a training set of question/answer pairs, and ...

Keywords: Web search, information retrieval, meta-search, query expansion, question answering

19 Consistency and replication: Consistency mechanisms for a distributed lookup service supporting mobile applications

Christoph Lindemann, Oliver P. Waldhorst

September 2003 **Proceedings of the 3rd ACM international workshop on Data engineering for wireless and mobile access MobiDe '03**

Publisher: ACM Press

Full text available: pdf(425.01 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents a general-purpose distributed lookup service, denoted Passive Distributed Indexing (PDI). PDI stores entries in form of (key, value) pairs in index caches located in each mobile device. Index caches are filled by epidemic dissemination of popular index entries. By exploiting node mobility, PDI can resolve most queries locally without sending messages outside the radio coverage of the inquiring node. Thus, PDI reduces network traffic for the resolution of keys to values. For k ...

Keywords: ad-hoc networked databases, consistency maintenance in weakly connected environments, data caching, epidemic algorithms for data dissemination

20 Special issue on Mobile Data Management: Exploiting epidemic data dissemination for consistent lookup operations in mobile applications

Christoph Lindemann, Oliver P. Waldhorst

July 2004 **ACM SIGMOBILE Mobile Computing and Communications Review**, Volume 8 Issue 3

Publisher: ACM Press

Full text available: pdf(391.34 KB) Additional Information: full citation, abstract, references

This paper presents a general-purpose distributed lookup service, denoted Passive Distributed Indexing (PDI). PDI stores entries in form of (key, value) pairs in index caches located at mobile devices. Index caches are filled by epidemic dissemination of popular index entries. By exploiting node mobility, PDI can resolve most queries locally without sending messages outside the radio coverage of the inquiring node. For keeping index caches coherent, configurable value timeouts implementing im ...

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

ACM PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
"search engine" and (timeout or "length of time") and "backg
THE ACM DIGITAL LIBRARY
Feedback Report a problem Satisfaction survey

Terms used: **search engine** and **timeout** or **length of time** and **background task**

Found **391** of **213,097**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 Cluster-based scalable network services
Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5
Publisher: ACM Press
Full text available: pdf(2.42 MB) Additional Information: full citation, references, citings, index terms

2 Aurora: a new model and architecture for data stream management
Daniel J. Abadi, Don Carney, Ugur Çetintemel, Mitch Cherniack, Christian Convey, Sangdon Lee, Michael Stonebraker, Nesime Tatbul, Stan Zdonik
August 2003 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 12 Issue 2
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(585.97 KB) Additional Information: full citation, abstract, citings, index terms

Abstract.This paper describes the basic processing model and architecture of Aurora, a new system to manage data streams for monitoring applications. Monitoring applications differ substantially from conventional business data processing. The fact that a software system must process and react to continual inputs from many sources (e.g., sensors) rather than from human operators requires one to rethink the fundamental architecture of a DBMS for this application area. In this paper, we present Aur ...

Keywords: Continuous queries, Data stream management, Database triggers, Quality-of-service, Real-time systems

3 Quantifying and Improving the Availability of High-Performance Cluster-Based Internet Services
Kiran Nagaraja, Neeraj Krishnan, Ricardo Bianchini, Richard P. Martin, Thu D. Nguyen
November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing SC '03**
Publisher: IEEE Computer Society
Full text available: pdf(306.01 KB) Additional Information: full citation, abstract

Cluster-based servers can substantially increase performance when nodes cooperate to

PORTAL
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

"search engine" and (timeout or "length of time") and (foreg

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used: **search engine** and **timeout** or **length of time** and **foreground** and **background task**

Found **18** of **213,097**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 18 of 18

Relevance scale

1 Visitor awareness in the web

Albrecht Schmidt, Hans-W. Gellersen

April 2001 **Proceedings of the 10th international conference on World Wide Web WWW '01**

Publisher: ACM Press

Full text available: pdf(863.11 KB) Additional Information: full citation, references, citings, index terms

Keywords: Web awareness, Web-based interaction, ambient media, augmented reality, visitor awareness, web activity

2 Idletime scheduling with preemption intervals

Lars Eggert, Joseph D. Touch

October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5

Publisher: ACM Press

Full text available: pdf(2.02 MB) Additional Information: full citation, abstract, references, index terms

This paper presents the idletime scheduler; a generic, kernel-level mechanism for using idle resource capacity in the background without slowing down concurrent foreground use. Many operating systems fail to support transparent background use and concurrent foreground performance can decrease by 50% or more. The idletime scheduler minimizes this interference by partially relaxing the work conservation principle during *preemption intervals*, during which it serves no background requests eve ...

Keywords: background processing, disk scheduler, idletime scheduling, network scheduler, preemption interval, queuing, resource scheduler

3 Cluster-based scalable network services

Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: pdf(2.42 MB) Additional Information: full citation, references, citings, index terms

4 Distributed systems: Dynamo: amazon's highly available key-value store

Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, Werner Vogels

October 2007 **Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles SOSP '07**

Publisher: ACM Press

Full text available: pdf(894.07 KB) Additional Information: full citation, abstract, references, index terms

Reliability at massive scale is one of the biggest challenges we face at Amazon.com, one of the largest e-commerce operations in the world; even the slightest outage has significant financial consequences and impacts customer trust. The Amazon.com platform, which provides services for many web sites worldwide, is implemented on top of an infrastructure of tens of thousands of servers and network components located in many datacenters around the world. At this scale, small and large components ...

Keywords: performance, reliability, scalability

5 A method for evaluating the impact of software configuration parameters on e-commerce sites

Monchai Sopitkamol, Daniel A. Menascé

July 2005 **Proceedings of the 5th international workshop on Software and performance WOSP '05**

Publisher: ACM Press

Full text available: pdf(276.37 KB) Additional Information: full citation, abstract, references

E-commerce systems are composed of many compon ents with several configurable software parameters which, if properly configured, can optimize system performance. The cost and burden of implementing upgrades can be deferred by judicial tuning of these parameters. Tuning is a time consuming effort and energy has to be spent on the most relevant parameters. This paper provides a method to evaluate the statistical significance of configurable parameters in e-commerce systems, the interaction effects ...

6 Semantics and dialog: Augmenting conversational dialogue by means of latent semantic googling

Robin Senior, Roel Vertegaal

October 2005 **Proceedings of the 7th international conference on Multimodal interfaces ICMI '05**

Publisher: ACM Press

Full text available: pdf(497.75 KB) Additional Information: full citation, abstract, references, index terms

This paper presents Latent *Semantic Googling*, a variant of Landauer's Latent Semantic Indexing that uses the Google search engine to judge the semantic closeness of sets of words and phrases. This concept is implemented via *Ambient Google*, a system for augmenting conversations through the classification of discussed topics. Ambient Google uses a speech recognition engine to generate Google keyphrase queries directly from conversations. These queries are used to analyze the semantic ...

Keywords: augmented intelligence, context, latent semantic indexing, speech

7 Aurora: a new model and architecture for data stream management

Daniel J. Abadi, Don Carney, Ugur Çetintemel, Mitch Cherniack, Christian Convey, Sangdon

Lee, Michael Stonebraker, Nesime Tatbul, Stan Zdonik
August 2003 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 12 Issue 2
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(585.97 KB) Additional Information: full citation, abstract, citings, index terms

Abstract.This paper describes the basic processing model and architecture of Aurora, a new system to manage data streams for monitoring applications. Monitoring applications differ substantially from conventional business data processing. The fact that a software system must process and react to continual inputs from many sources (e.g., sensors) rather than from human operators requires one to rethink the fundamental architecture of a DBMS for this application area. In this paper, we present Aur ...

Keywords: Continuous queries, Data stream management, Database triggers, Quality-of-service, Real-time systems

8 E-commerce: Ranking configuration parameters in multi-tiered e-commerce sites

Monchai Sopitkamol
December 2004 **ACM SIGMETRICS Performance Evaluation Review**, Volume 32 Issue 3
Publisher: ACM Press
Full text available: pdf(782.16 KB) Additional Information: full citation, abstract, references

E-commerce systems are composed of many compon ents with several configurable parameters that, if properly configured, can optimize system performance. Before upgrading existing systems to overcome performance bottlenecks, several areas of a site's architecture and its parameters may be adjusted to improve performance. This paper provides a method to rank key configurable e-commerce system parameters that significantly impact overall system performance, and the performance of the most significant ...

9 Quantifying and Improving the Availability of High-Performance Cluster-Based Internet Services

Kiran Nagaraja, Neeraj Krishnan, Ricardo Bianchini, Richard P. Martin, Thu D. Nguyen
November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing SC '03**
Publisher: IEEE Computer Society
Full text available: pdf(306.01 KB) Additional Information: full citation, abstract

Cluster-based servers can substantially increase performance when nodes cooperate to globally manage resources. However, in this paper we show that cooperation results in a substantial availability loss, in the absence of high-availability mechanisms. Specifically, we show that a sophisticated cluster-based Web server, which gains a factor of 3 in performance through cooperation, increases service unavailability by a factor of 10 over a non-cooperative version. We then show how to augment this W ...

10 DDoS defense by offense

Michael Walfish, Mythili Vutukuru, Hari Balakrishnan, David Karger, Scott Shenker
August 2006 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '06**, Volume 36 Issue 4
Publisher: ACM Press
Full text available: pdf(334.96 KB) Additional Information: full citation, abstract, references, index terms

This paper presents the design, implementation, analysis, and experimental evaluation of *speak-up*, a defense against *application-level* distributed denial-of-service (DDoS), in which attackers cripple a server by sending legitimate-looking requests that consume

computational resources (*e.g.*, CPU cycles, disk). With speak-up, a victimized server encourages all clients, resources permitting, *to automatically send higher volumes of traffic*. We suppose that attackers are a ...

Keywords: DoS attack, bandwidth, currency

11 Thread-level parallelism and interactive performance of desktop applications

Kristián Flautner, Rich Uhlig, Steve Reinhardt, Trevor Mudge

November 2000 **ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , Proceedings of the ninth international conference on Architectural support for programming languages and operating systems ASPLOS-IX**, Volume 28 , 34 Issue 5 , 5

Publisher: ACM Press

Full text available: pdf(234.58 KB) Additional Information: full citation, abstract, references, citings, index terms

Multiprocessing is already prevalent in servers where multiple clients present an obvious source of thread-level parallelism. However, the case for multiprocessing is less clear for desktop applications. Nevertheless, architects are designing processors that count on the availability of thread-level parallelism. Unlike server workloads, the primary requirement of interactive applications is to respond to user events under human perception bounds rather than to maximize end-to-end throughput. In ...

12 Thread-level parallelism and interactive performance of desktop applications

Krisztián Flautner, Rich Uhlig, Steve Reinhardt, Trevor Mudge

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Publisher: ACM Press

Full text available: pdf(2.94 MB) Additional Information: full citation, abstract, references, citings, index terms

Multiprocessing is already prevalent in servers where multiple clients present an obvious source of thread-level parallelism. However, the case for multiprocessing is less clear for desktop applications. Nevertheless, architects are designing processors that count on the availability of thread-level parallelism. Unlike server workloads, the primary requirement of interactive applications is to respond to user events under human perception bounds rather than to maximize end-to-end throughput. In ...

13 A case study of automatically creating test suites from web application field data

Sara Sprenkle, Emily Gibson, Sreedevi Sampath, Lori Pollock

July 2006 **Proceedings of the 2006 workshop on Testing, analysis, and verification of web services and applications TAV-WEB '06**

Publisher: ACM Press

Full text available: pdf(184.56 KB) Additional Information: full citation, abstract, references, index terms

Creating effective test cases is a difficult problem, especially for web applications. To comprehensively test a web application's functionality, test cases must test complex application state dependencies and concurrent user interactions. Rather than creating test cases manually or from a static model, field data provides an inexpensive alternative to creating such sophisticated test cases. An existing approach to using field data in testing web applications is user-session-based testing. Previ ...

14 MusicFX: an arbiter of group preferences for computer supported collaborative workouts

Joseph F. McCarthy, Theodore D. Anagnost

November 1998 **Proceedings of the 1998 ACM conference on Computer supported cooperative work CSCW '98**

Publisher: ACM Press
Full text available: pdf(1.14 MB) Additional Information: full citation, references, citings, index terms

Keywords: audio, audiospaces, empirical studies, evaluation, human-computer interaction, intelligent environments, shared spaces, social behavior, social interactions, ubiquitous computing

15 Design and test strategies for a safety-critical embedded executive

Charles A. Meyer, Michael G. Reznick
December 1996 **Proceedings of the conference on TRI-Ada '96: disciplined software development with Ada TRI-Ada '96**
Publisher: ACM Press
Full text available: pdf(900.36 KB) Additional Information: full citation, references, index terms

16 Scalability in MMOGs: Load balancing for massively multiplayer online games

Fengyun Lu, Simon Parkin, Graham Morgan
October 2006 **Proceedings of 5th ACM SIGCOMM workshop on Network and system support for games NetGames '06**
Publisher: ACM Press
Full text available: pdf(544.67 KB) Additional Information: full citation, abstract, references, index terms

Supporting thousands, possibly hundreds of thousands, of players is a requirement that must be satisfied when delivering server based online gaming as a commercial concern. Such a requirement may be satisfied by utilising the cumulative processing resources afforded by a cluster of servers. Clustering of servers allow great flexibility, as the game provider may add servers to satisfy an increase in processing demands, more players, or remove servers for routine maintenance or upgrading. If ca ...

17 Charting past, present, and future research in ubiquitous computing

Gregory D. Abowd, Elizabeth D. Mynatt
March 2000 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 7 Issue 1
Publisher: ACM Press
Full text available: pdf(730.83 KB) Additional Information: full citation, abstract, references, citings, index terms

The proliferation of computing into the physical world promises more than the ubiquitous availability of computing infrastructure; it suggest new paradigms of interaction inspired by constant access to information and computational capabilities. For the past decade, application-driven research on abiquitous computing (ubicomp) has pushed three interaction themes:natural interfaces, context-aware applications,andautomated capture and access. To chart a cours ...

Keywords: augmented reality, capture and access, context-aware applications, evaluation, everyday computing, natural interfaces, social implications, ubiquitous computing, user interfaces

18 Tradeoffs in displaying peripheral information

Paul P. Maglio, Christopher S. Campbell
April 2000 **Proceedings of the SIGCHI conference on Human factors in computing systems CHI '00**
Publisher: ACM Press

Full text available: pdf(939.18 KB) Additional Information: full citation, abstract, references, citings, index terms

Peripheral information is information that is not central to a person's current task, but provides the person the opportunity to learn more, to do a better job, or to keep track of less important tasks. Though peripheral information displays are ubiquitous, they have been rarely studied. For computer users, a common peripheral display is a scrolling text display that provides announcements, sports scores, stock prices, or other news. In this paper, we investigate how to design peripheral disp ...

Keywords: dual-task tradeoffs, peripheral information, user interface design

Results 1 - 18 of 18

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

Fri, 26 Oct 2007, 11:04:02 AM EST

Edit an existing query or compose a new query in the Search Query Display.

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Search Query Display

Run Search Reset

Recent Search Queries

#1 (background or foregruond processing) <in> metadata

#2 (background or foregruond processing) <in> metadata

#3 (background or foreground) and task or search

#4 (background search and (timeout)<IN>metadata)

#5 (background search and (timeout)<IN>metadata)

#6 (~~background serach~~ and timeout<IN>metadata)

#7 (~~background serach~~ and timeout<IN>metadata)

#8 (~~background serach~~ and timeout<IN>metadata)

#9 (backgound and task<IN>metadata)

#10 (backgound and task<IN>metadata)

#11 (backgound and task and searching and timeout<IN>metadata)

#12 (backgound and task and searching and timeout<IN>metadata)

#13 (background and task and search and timeout<IN>metadata)

Clear Session History

© Copyright 2006 IEEE –

Indexed by Inspec